

Denial-of-Service Attacks in Wireless Networks using off-the-shelf Hardware

Alexandros Fragkiadakis*, Ioannis Askoxylakis, Panos Chatziadam

Institute of Computer Science, Foundation for Research and Technology - Hellas
(FORTH)

P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece
{alfrag,asko,panosc}@ics.forth.gr

Abstract. Wireless network technologies offer ubiquitous broadband access to millions of users at an affordable cost. However, the broadband nature of the wireless medium make these networks vulnerable to a number of attacks. Malicious interference at the physical layer, and extended packet collisions at the medium access layer can cause significant DoS attacks. In this work, we show how off-the-shelf hardware can be used to create devastating DoS attacks in a IEEE 802.11 network. Moreover, we present two algorithms for attack detection that are based on the cumulative sum algorithm.

1 Introduction

Network proliferation has been remarkable, especially during the last decade. Technology advancements in the area of network communications have offered high performance improvement and ubiquitous Internet access. From the era of the early-stage communication protocols (e.g. Aloha) through the active networks [1], and software-defined networks [2], ubiquitous network access has been achieved thanks to the advances in wireless technologies. A number of communication protocols (IEEE 802.11, IEEE 802.15.4, IEEE 802.16, etc.) enable energy efficient communications in relatively large distances. Thousands of smart phones and other wireless clients can now enjoy any-time any-where Internet connectivity.

Nevertheless, the broadcast nature of the wireless medium make wireless communications susceptible to a number of threats. Adversaries can cause severe Denial-of-Service attacks (DoS) by exploiting a number of vulnerabilities. DoS attacks pose a major threat in every communication system, often with catastrophic results as wireless communications are nowadays used in many applications (e.g. smart cards [3]). At the physical layer, adversaries can generate malicious interference resulting in heavy packet loss in the network. At the medium access (MAC) layer, malicious users can create extended packet collisions, causing severe DoS attacks in the wireless network. Throughout this work we refer to adversaries and malicious users as jammers.

* Corresponding author

Our contribution focuses on showing how off-the-shelf equipment can be used to create DoS attacks at the physical, and medium access layers. We also present two algorithms for the detection of these attacks. The rest of this paper is organized as follows. In Section 2 we present the basic components of our off-the-shelf jammer. Section 3 describes how malicious interference can be used to create DoS attacks. In Section 4 we demonstrate how extended packet collisions are easily caused by jammers, and how greedy behaviors affect network’s performance. Section 5 presents two algorithms for attack detection based on the cumulative-sum algorithm. Finally, conclusions appear in Section 6.

2 Off-the-shelf hardware for malicious purposes

As mentioned in the previous section, wireless networks, due to their broadcast nature, are susceptible to a number of threats. A major threat referred as physical-layer jamming refers to interference created by a malicious node. There is a number of commercial devices that can be used for this purpose (e.g. [4–6]). Nevertheless, as we show in this paper, off-the-shelf hardware can be successfully used to launch severe DoS attacks in a wireless network.

Our jamming equipment is based on a mini-ITX board (Fig. 1) carrying 512 MB of RAM with a 80 GB disk. The wireless interface cards are based on the Atheros 802.11a/b/g CM9-GP mini-PCI card. Furthermore, Ath5k [7], an open source IEEE 802.11 driver is used, on Gentoo Linux.



Fig. 1: Off-the-shelf jamming device

The software part of the jammer is shown in Fig. 2. This consists of several components implemented in both the kernel and user spaces of the Linux operating system. At user-space, the *command repository* contains all the attack characteristics. These define a detailed adversary model to be used against a wireless network. Such a typical model can describe, for example, the wireless channel to be attacked, the attack intensity in terms of packet rate, transmission power, attack duration, etc. Commands are propagated through the *netlink socket interface* to kernel-space, stored in the *command trace collection* module

that provides buffering capabilities. Finally, the characteristics of the adversary model are used to setup several parameters of the Ath5k driver in order to make an attack feasible.

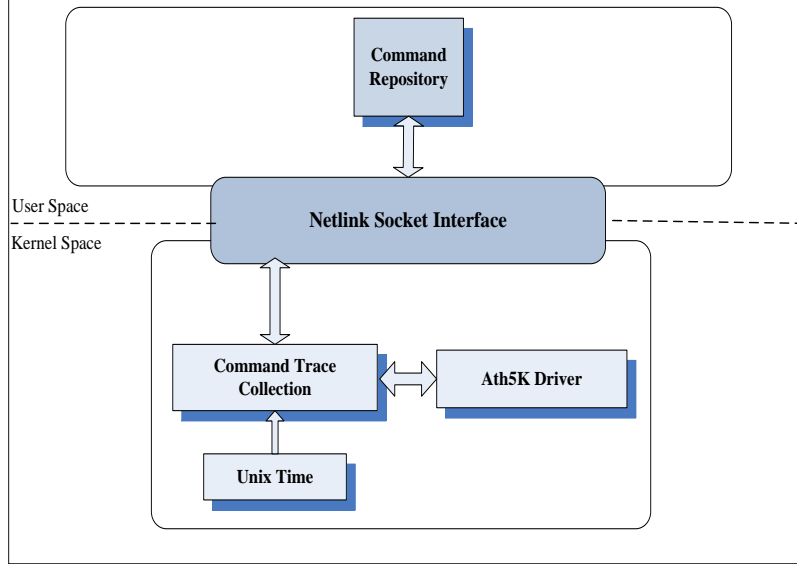


Fig. 2: Jammer software layout

3 Physical-layer attacks

A major threat in wireless networks is interference that is caused by signal emissions in neighboring channels. In general, interference can be characterized as malicious or non-malicious, depending on the incentives of the interferer. Non-malicious interference can be caused by nearby legitimate nodes that operate in neighboring channels ([8]). Malicious physical-layer interference (jamming) is created by signal emissions in neighboring channels. This affects both the transmitters and the receivers of a wireless network. IEEE 802.11 transmitters sense the wireless medium before any transmission takes place. If the measured noise is above a threshold, they refrain from transmission for some random time. So, if a jammer is present, the transmission operations of the legitimate nodes can be heavily disrupted, hence DoS attacks become feasible. On the other hand, legitimate receivers cannot correctly detect and decode incoming packets in the presence of jammer. This is due to the excessive noise generated in their vicinity that leads to an extensive packet loss. Moreover, as packets are lost in the network, further retransmissions by the transmitters take place causing severe network disruption.

In order to demonstrate the effects of physical-layer jamming we use a single off-the-shelf jammer with characteristics as described in Section 2. The specific type of jammer does not follow any rules of the IEEE 802.11 protocol, so it freely performs jamming even if legitimate transmissions are taking place. We setup a network of four legitimate nodes: Sender, Receiver, Monitor1, and Monitor2. Packets flow from Sender to Receiver, while periodic jamming using the frequency of the neighboring channel is taking place. Receiver, Monitor1, and Monitor2 record the packets that belong to the legitimate traffic, and for every recorded packet, the SINR (Signal-to-Interference plus Noise Ratio) is computed, taking into account the power leakage in the neighboring channels as described in [9]. Fig. 3 shows how SINR substantially drops during the jamming attacks (symbolized by the orthogonal boxes at the bottom of each graph).




Fig. 3: SINR variations during the jamming attacks

When SINR significantly drops, a wireless receiver often becomes unable to detect and further decode a transmitted packet. This results to packet loss, throughput degradation, as well as energy waste, as transmitters keep re-transmitting packets. Fig. 4 shows the retry attempts of the transmitter (Sender), the throughput at the Receiver, and the total packet loss in the network when a jammer is present. Retry attempts increase up to five times during jamming, while throughput drops to about 1 Mbps (from 15 Mbps when no jamming takes place), and packet loss increases over 60%.




Fig. 4: Performance degradation during the jamming attacks

4 MAC layer attacks

IEEE 802.11 is a CSMA/CA (carrier sense multiple access with collision avoidance) based medium access protocol. Each potential transmitter has to first sense the wireless medium, and if it is free, transmission takes place. However, if the medium is occupied, it has to enter a back-off stage where it waits for some amount of time before repeating the same procedure (sensing, etc) [10].

4.1 Denial-of-service attacks through packet collisions

Attackers can exploit the CSMA/CA mechanism of IEEE 802.11 by emitting energy when the wireless medium is occupied by a legitimate node. At this point, jammer emits energy on the same channel legitimate nodes use for communication, aiming to cause packet collisions and to degrade network's performance. Packet collisions refer to captured packets that mainly suffer from CRC (cyclic redundancy check) errors. Fig. 5 shows the ratio of the corrupted packets (CRC errors) over the correctly decoded ones, captured in a single receiver when a periodic jammer is present. Jammer operates on the same channel used for the legitimate communication. Observe that the ratio exceeds 60% when jamming traffic is emitted. This is a severe DoS attack as corrupted packets are essentially lost packets that the sender will attempt to re-transmit up to a number of times (retry limit).

Next, we demonstrate how jamming on the same channel affects the performance of video transmission in the wireless network. For this reason, we set up




Fig. 5: Ratio of corrupted packets over the correctly decoded ones

a testbed consisting of a video server, a video client and the jammer described in Section 2. Encoded MPEG-4 video is transmitted from server to client using RTP/UDP packets. Periodic jamming takes place in the network. We measure video's performance using PSNR (Peak-Signal-to-Noise Ratio), an objective quality metric widely used to measure video performance. Supposing there are two $m \times n$ images S and D , where S is the original image and D the reconstructed image, the PSNR of this image is given by:

$$PSNR = 20 \times \log_{10} \frac{V_{peak}}{\sqrt{MSE}} \quad (1)$$

where V_{peak} is its maximum value (e.g. 255 for 8-bit encoding), and MSE is the mean squared error given by:

$$MSE = \frac{1}{m \times n} \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} [S(i, j) - D(i, j)]^2 \quad (2)$$

Fig. 6 [11] shows the PSNR for each received video frame, and for two experiments: (i) when no jamming is used (No Jam), and (ii) when jammer is active periodically for a duration of 80 seconds (Jam). Observe that PSNR significantly drops when jamming is taking place.

4.2 Greedy behavior

The CSMA/CA mechanism of IEEE 802.11 requires that potential transmitters should wait for some time when the wireless medium is busy in order to decrease the collision probability. The waiting time (back-off time) is chosen uniformly




Fig. 6: PSNR per frame

in the interval $[0 - CW]$, where CW is the *contention window* size. Initially, CW equals CW_{min} , that is the *minimum contention window*. Each time a node finds the wireless medium busy, it doubles CW up to CW_{max} . When a successful transmission takes place, CW reduces to CW_{min} .

This mechanism can be exploited by a malicious (or greedy) node assigning a very small value to its CW_{min} . With a small CW_{min} , the malicious node can monopolize the medium and make the legitimate nodes entering the back-off stage repeatedly. Fig. 7 shows the throughput achieved by four wireless nodes when all attempt to transmit a UDP flow of 200 Kbytes to a single access point, without the presence of a greedy node.

Next, Node1 becomes greedy by periodically assigning a very small value to its CW_{min} . Repeating the same experiment, in Fig. 8 we show that Node1's throughput increases from about 200 Kbytes/sec to 350 Kbytes/sec while the throughput of the rest of the nodes falls almost to 100 Kbytes/sec. Node1 becomes greedy every 10 seconds for a duration of 10 seconds. After that period, it stops behaving greedy by assigning a proper value to its CW_{min} .

Such greedy behaviors negatively affect legitimate nodes performance reducing the fairness of the wireless system. Fairness is related to the ability of the MAC layer to equitably share a common channel between a number of contending nodes [12]. Jain's fairness index [13] is widely used as a metric to measure the fairness of a system. Assuming that N is the number of competing flows and γ_i the fraction of packets from node i that arrived within a time window, Jain's index is defined as follows:




Fig. 7: Throughput with the absence of a greedy node





Fig. 8: Throughput with the presence of a greedy node (Node1)


$$F = \frac{(\sum_{i=1}^N \gamma_i)^2}{N \times \sum_{i=1}^N \gamma_i^2} \quad (3)$$

When $F = 1$ perfect fairness is achieved, while when $F = \frac{1}{N}$ absolute unfairness is achieved.

Fig. 9a shows how the fairness of the network drops when one of the nodes becomes greedy. Observe that as the number of flows decreases, the drop in fairness increases. This is because less nodes content for the medium, hence it is easier for the greedy node to monopolize it by selecting a small CWmin value. On the other hand, if no greedy node is present, fairness increases, for all flows (Fig. 9b).



(a) When a greedy node is present



(b) When a greedy node is absent

Fig. 9: Fairness for a different number of flows

5 Attack detection

In this section we describe techniques for the detection of physical-layer jamming, and collisions at the MAC layer. As shown in Fig. 3, jamming causes extended SINR drops. Based on this, we deploy a cumulative-sum (cusum) algorithm [14] able to detect abrupt changes of the SINR. In previous works [15, 11, 16, 17] we show that maximum performance, in terms of false alarms/detection probability, is achieved when considering the maximum minus the minimum values of SINR within a short and long windows. Cusum is defined as:

$$y_n = \begin{cases} y_{n-1} + Z_n - a & \text{if } y_n \geq 0 \\ 0 & \text{if } y_n < 0 \end{cases} \quad (4)$$

Z_n is the expectation of a specific metric that changes whenever jamming takes place (in our case the SINR-based metric), and $a \in R^+$ controls its drift. Furthermore, Z_n is given by $Z_n = D(n) - \bar{D}(n)$, where

$$D(n) = \max_{n-K+1 < i \leq n} x_i - \min_{n-K+1 < i \leq n} x_i,$$

and

$$\bar{D}(n) = \frac{\sum_{i=n-M+1}^n D(i)}{M},$$

During the jamming attacks, cusum's output increases as shown in Fig. 10, and when it exceeds a predefined threshold, an alarm is raised.




Fig. 10: Cusum's output

We use the same technique for the detection of MAC layer attacks in cases where an adversary causes packet collisions in the network. Rather than using

the SINR as metric for the cusum, we consider the ratio of the corrupted packets over the correctly decoded ones (Fig. 4). Cusum's output increases similarly as in Fig. 10 indicating the attack.

6 Conclusions

In this paper we demonstrated how commodity hardware can be used as jammer and severely affect network's performance. At the physical layer, jammer can create malicious interference that substantially degrades network performance. At the MAC layer, jammer can create extended packet collisions by energy emission when legitimate wireless traffic is being transmitted. Furthermore, a greedy node by exploiting the back-off mechanism of IEEE 802.11, can monopolize the medium, restricting network resources to the non-misbehaving nodes.

Efficient detection of these attacks is feasible by considering an SINR-based metric at the physical layer, and the ratio of the corrupted packets over the correctly decoded ones at the MAC layer. Both metrics are utilized by a cusum algorithm that signals an alarm if a predefined threshold has been exceeded.

References

1. N. Bartzoudis, A. Fragkiadakis, D. Parish, J. Luis-Nunez, and M. Sandford, "Re-configurable computing and active networks," in *Engineering of Reconfigurable Systems and Algorithms*, 2003, pp. 280–283.
2. C. Monsanto, J. Reich, N. Foster, J. Rexford, and D. Walker, "Composing software-defined networks," in *Proc. of the NSDI*, 2013, pp. 87–98.
3. K. Markantonakis, K. Mayes, D. Sauveron, and I. Askoxylakis, "Overview of security threats for smart cards in the public transport industry," in *Proc. of the ICEBE*, 2008, pp. 506–513.
4. "Sesp jammers, <http://www.sesp.com>."
5. "Mobile device jammer, <http://www.phonejammer.com/home.php>."
6. "Software-defined radios, <http://www.ettus.com/home>."
7. "Linux wireless drivers, ath5k, <http://linuxwireless.org/en/users/Drivers/ath5k>."
8. E. Tragos, A. Fragkiadakis, I. Askoxylakis, and V. Siris, "The impact of interference on the performance of a multi-path metropolitan wireless mesh network," in *Proc. of ISCC*, 2011, pp. 199–204.
9. V. Angelakis, S. Papadakis, V. Siris, and A. Traganitis, "Channel Interference in 802.11a is harmful. Testbed validation of a simple quantification model," *IEEE Communications Magazine*, pp. 160–166, 2011.
10. M. Natkaniec and A. Pach, "An analysis of the backoff mechanism used in ieee 802.11 networks," in *Proc. of ISCC*, 2000, pp. 444–449.
11. A. Fragkiadakis, E. Tragos, and I. Askoxylakis, "Video streaming performance in wireless hostile environments," in *Proc. of the 5th FTRA International Conference on Multimedia and Ubiquitous Engineering*, 2011, pp. 267–272.
12. G. Berger-Sabbatel, A. Duda, M. Heusse, and F. Rousseau, "Short-term fairness of 802.11 networks with several hosts," in *Proc. of the IFIP*, 2004, pp. 263–274.
13. R. Jain, *The Art of Computer Systems Performance Analysis*. John Wiley & Sons, 1991.

14. A. Cardenas, S. Radosavac, and J. Baras, "Evaluation of detection algorithms for mac layer misbehavior: Theory and experiments," *IEEE/ACM Transactions on Networking*, 2009.
15. A. Fragkiadakis, V. Siris, and A. Traganitis, "Effective and robust detection of jamming attacks," in *Proc. of Future Network and Mobile Summit*, 2010, pp. 1–8.
16. A. Fragkiadakis, E. Tragos, and I. Askoxylakis, "Design and Performance Evaluation of a Lightweight Wireless Early Warning Intrusion Detection Prototype," *EURASIP Journal on Wireless Communications and Networking*, vol. 12, pp. 1–18, 2012.
17. A. Fragkiadakis, V. Siris, N. Petroulakis, and A. Traganitis, "Anomaly-based intrusion detection of jamming attacks, local versus collaborative detection," *Wireless Communications and Mobile Computing*, pp. 1–19, 2013.